Centre de recherche sur les transports (C.R.T.)
Centre for Research on Transportation

CRT-2004-12

THE MAXIMAL EXPECTED COVERAGE RELOCATION PROBLEM FOR EMERGENCY VEHICLES

by

Michel Gendreau
Gilbert Laporte
Frédéric Semet

May 2004

Université de Montréal
C.P. 6128, succursale Centre-ville, Montréal QC H3C 3J7 Canada
Téléphone: (514) 343-7575 / Télécopieur: (514) 343-7121

crt@crt.umontreal.ca
www.crt.umontreal.ca

HEC ÉCOLE POLYTECHNIQUE MONTRÉAL Université de Montréal

PUBLICATION
CRT-2004-12

THE MAXIMAL EXPECTED COVERAGE RELOCATION PROBLEM FOR EMERGENCY VEHICLES

by

Michel Gendreau[1]
Gilbert Laporte[2]
Frédéric Semet[3]

The research was supported in part by the Fonds de la recherche en santé du Québec (FRSQ), and by the Natural Sciences and Engineering Research Council of Canada (NSERC).

1. Département d'informatique et de recherche opérationnelle and Centre de recherche sur les transports, Université de Montréal, C.P. 6128, succursale Centre-ville, Montréal, Canada H3C 3J7

2. Chaire de recherche du Canada en distributique, HEC Montréal, 3000 Côte-Sainte-Catherine, Montréal, Canada H3T 2A7, and Centre de recherche sur les transports, Université de Montréal, C.P. 6128, succursale Centre-ville, Montréal, Canada H3C 3J7

3. LAMIH-ROI, Université de Valenciennes et du Hainaut-Cambrésis, Le Mont-Houy, 59313 Valenciennes Cedex 9, France and GRIS, Département d'administration de la santé, Université de Montréal, C.P. 6128, succursale Centre-ville, Montréal, Canada H3C 3J7

Département d'informatique et de recherche opérationnelle - Publication #1248

Dépôt légal – Bibliothèque nationale du Québec
Bibliothèque nationale du Canada, 2004

Abstract

In the *Maximal Expected Coverage Relocation Problem* the aim is to provide a dynamic relocation strategy for emergency vehicles waiting sites in such a way that the expected covered demand is maximized and the number of waiting site relocations is controlled. The problem can be formulated as an integer linear program. When the number of vehicles is relatively small this program can be solved within reasonable computing time. Simulations conducted with real-life emergency medical services data from the Montreal area confirm the feasibility of the proposed approach.

Keywords : location, relocation, emergency vehicles, physician cars, ambulances, coverage models, dynamic models.

Résumé

Le problème de relocalisation avec maximisation de la couverture espérée a pour but de proposer une stratégie de repositionnement des points d'attente des véhicules d'urgence de façon à maximiser la demande couverte espérée, le nombre de relocalisation de sites d'attente étant borné. Le problème peut être formulé sous la forme d'un programme linéaire en nombres entiers. Lorsque le nombre de véhicules n'est pas trop élevé, ce programme peut être résolu en un temps de calcul raisonnable. Des simulations réalisées sur des données fournies par le service de véhicules d'urgence de la région de Montréal confirment la faisabilité de l'approche proposée.

Mots-clefs : localisation, relocalisation, véhicules d'urgence, véhicules de médecin, ambulances, modèles de couverture, modèles dynamiques.

Introduction

In most major cities, emergency medical services are provided by a fleet of ambulances complemented by physician cars. These are dispatched in addition to an ambulance to a limited number of high-emergency calls requiring the presence of a physician, which include car accidents involving major injuries, some cardiac arrest cases, etc. The physician vehicles must be located at appropriate points in order to provide adequate coverage. Whenever a vehicle is dispatched to a call, it may leave a significant fraction of the population without coverage; it is therefore customary for ambulance services to apply a relocation policy in order to maintain a proper coverage level. Our study is motivated by the problem of locating and relocating physician cars in the Montreal area. It provides a planning tool to optimally relocate these vehicles. While the number of ambulances serving the Montreal area (population : 2.5 million) varies between 40 and 60 depending on the time of the day, only a few physician cars are needed, typically between three and six. As will be seen, this small number makes our approach feasible

Most available studies on such problems apply to ambulances only and very few address the relocation issue. There has been a marked evolution over the past three decades in the development of ambulance location models and algorithms. Some of the earliest models are the *Location Set Covering Problem* (LSCP) proposed by Toregas *et al.*[1] and the *Maximal Covering Location Problem* (MCLP) introduced by Church and ReVelle[2]. In the LSCP, the aim is to minimize the number of vehicles required to cover a set of demand points. In this model, a point is said to be *covered* if it can be reached by at least one vehicle within a preset driving time r. The MCLP works with a given number n of vehicles and attempts to cover the largest possible demand. Both models make sense in practice. The first can be used to determine an appropriate fleet size to meet a certain demand whereas the second can help optimally allocate available but insufficient resources. A good compromise between coverage and cost can be reached by solving the MCLM with various values of n and selecting one of the solutions. Several extensions of these two basic models have been proposed by Schilling *et al.*[3], Daskin and Stern[4], Hogan and ReVelle[5], and Gendreau, Laporte and Semet[6].

In practice such *static models* are inadequate in the sense that their solution is unlikely to provide sufficient coverage once a vehicle is dispatched to a call. One way around this difficulty has been

to develop *probabilistic models* which reflect the fact that emergency vehicles are busy only a fraction of the time and become unavailable once they respond to a call. They must in fact be considered as servers that operate within a queueing system. This line of research was pioneered by Larson[7,8,9]. Well-known probabilistic covering models include the *Maximum Expected Covering Problem* (MEXCLP) proposed by Daskin[10], its extensions TIMEXCLP[11], AMEXCLP[12] and QPLSCP[13], the *Maximal Availability Location Problems* (MALP I and MALP II) of ReVelle and Hogan[14], and the *Rel-P Model* of Ball and Lin[15]. These models make different assumptions on travel time distribution, on the spatial dependency of vehicle availability, and on desired probabilistic covering thresholds. A third approach is the use of *dynamic models*. Here, instead of seeking a single solution to a static or probabilistic model, the idea is to dynamically relocate vehicles in real-time as vehicles are dispatched to calls. Each relocation amounts to solving a static model subject to side constraints on vehicle moves. For example, one should avoid relocating too many vehicles at once or moving the same vehicle too often over a short period. An early dynamic model was proposed by Kolesar and Walker[16] for the relocation of fire companies. More recently Gendreau, Laporte and Semet[17] have developed a dynamic ambulance relocation model which can be applied in real-time through the use of parallel computing. For surveys of emergency vehicle location and relocation problems, see Marianov and ReVelle[18] and Brotcorne, Laporte and Semet[19].

One drawback of dynamic relocation algorithms is the need to compute a new solution whenever a vehicle is dispatched to a call. This can be time consuming or even infeasible when calls arrive in quick succession throughout the day, as was the case in the Island of Montreal ambulance relocation problem studied by Gendreau, Laporte and Semet[17]. Since parallel computing is not always a practical option in emergency vehicle dispatching centers, an alternative is to precompute a series of location scenarios for $k = 1,\ldots, n$ vehicles which can readily be applied whenever a call is made. This approach will work as long as n is not too large, as is the case in our application.

Our aim is to formulate and solve a dynamic problem arising in the relocation of physician vehicles. *Events* corresponding to emergency calls occur randomly at discrete time instants during the day. Whenever an event occurs, a vehicle is dispatched to the call location and a fleet relocation may take place. Each solution maximizes coverage given the number of available

vehicles. The objective is to maximize the expected coverage over time subject to an upper bound on the number of waiting sites that can be changed at each event. This problem is called the *Maximal Expected Coverage Relocation Problem* (MECRP). The MECRP will be formulated in the next section, followed by computational results on real-life data, and by some conclusions.

Mathematical model

There are two main ways of solving dynamic optimization problems. The *a posteriori* approach is the most common : it consists of computing and implementing a new solution whenever an event occurs. This is the strategy used in several dynamic vehicle routing algorithms (see, *e.g.*, Gendreau *et al.*[20]). The main drawback of this approach is that there is not always sufficient time to recompute a solution between two successive events if these occur in quick succession. To alleviate this difficulty, Gendreau, Laporte and Semet[17] have proposed an *a priori* methodology in which several solutions are precomputed in anticipation of future events and the appropriate solution, if available, is implemented whenever an event occurs. This methodology was successfully applied to an ambulance relocation problem in Montreal. It was shown that a precomputed solution was available in more than 95% of events. In the remaining cases the decision was to perform no vehicle relocation and simply wait for the next event to occur. The approach we develop in this study goes further in this direction. It precomputes a *unique* solution at the beginning of the planning period. This solution provides a full list of waiting sites for each possible event that may occur. What makes this strategy feasible is that the number of system states is relatively small in the application under study.

The MECRP is defined on a directed graph $G = (V \cup W, A)$ where V is the vertex set of demand points and W is a vertex set of potential waiting sites for $n \leq |W|$ emergency vehicles and A is a set of arcs defined on $(V \cup W)^2$. To each arc $(i, j) \in A$ is associated a driving time. Let d_i be the demand at vertex $i \in V$. As is often the case in location, it is assumed that the problem is discretized so that d_i may correspond to the population of a zone associated with i, for example a city block. A vertex $i \in V$ is said to be covered by a vertex $j \in W$ whenever the driving time from j to i is less than a prespecified coverage radius r. Denote by W_i the subset of vertices of W covering $i \in V$.

It is useful at this stage to recall the classical MCLP formulation for n vehicles. Let x_j be a binary variable equal to 1 if and only if a vehicle is located at $j \in W$ and let y_i be a binary variable equal to 1 if and only if demand point $i \in V$ is covered. Denote the solution value by z_n. The MCLP for n vehicles is then :

(MCLP-n) Maximize $z_n = \sum_{i \in V} d_i \, y_i$ (1)

subject to : $\sum_{j \in W_i} x_j \geq y_i$ $(i \in V)$ (2)

$\sum_{j \in W} x_j = n$ (3)

$x_j \in \{0,1\}$ $(j \in W)$ (4)

$y_i \in \{0,1\}$ $(i \in V)$. (5)

In this model, constraints (2) mean that vertex $i \in V$ is covered only if a least one vehicle is located in W_i, and constraints (3) control the number of vehicles used in the solution. A simple relaxation procedure for the MECRP consists of solving MCLP-k for $k = 1,\ldots, n$, and of applying the appropriate solution at each event. The expected solution cost is then $\sum_{k=1}^{n} q_k \, z_k$, where q_k is the probability of having k available vehicles. If the arrival rate of calls is λ and the average service rate is μ, then the probability that a vehicle is available is $p = 1 - \lambda/n\mu$, where $\lambda/n\mu$ corresponds to the busy fraction in MEXCLP[10] and related models. The value of q_k is easily obtained by means of a binomial distribution : $q_k = \binom{n}{k} p^k (1-p)^{n-k}$ $(k = 0,\ldots, n)$.

This relaxation procedure ignores constraints on the number of allowed waiting site changes at each event. To incorporate such constraints, it is useful to view the system as being in a succession of states k over time, where k is the number of available vehicles waiting for calls $(k = 0,\ldots, n)$. A given state k is described by binary variables x_{jk} equal to 1 if and only if a vehicle is located at $j \in W$, and by binary variables y_{ik} equal to 1 if and only if demand point $i \in V$ is covered by at least one vehicle. We assume that at most one vehicle is dispatched to a call at any event (if the system is in state 0, then obviously no vehicle can be dispatched). When the system

moves from state k to state $k + 1$ (if $k \leq n - 1$), at most α_k of the current k waiting sites can be changed; in the same vein, when the system moves from state k to state $k - 1$ (if $k \geq 1$), at most $\alpha_{k-1} + 1$ of the current waiting sites can be modified. This follows from the fact that α_{k-1} sites of configuration $k - 1$ are possibly modified and another site is discarded when going down from state k to state $k - 1$. To avoid multiple counting it is sufficient to keep track of waiting site changes when the system moves from state k to state $k + 1$ ($k = 1,\ldots, n - 1$). Hence define binary variables u_{jk} equal to 1 if and only if $j \in W$ ceases to be a waiting site when the system moves from state k to state $k + 1$. The MECRP is then :

$$\text{(MECRP)} \qquad \text{Maximize} \sum_{k=1}^{n} \sum_{i \in V} d_i \, q_k \, y_{ik} \qquad (6)$$

$$\text{subject to :} \quad \sum_{j \in W_i} x_{jk} \geq y_{ik} \qquad (i \in V, k = 0,\ldots, n) \qquad (7)$$

$$\sum_{j \in W} x_{jk} = k \qquad (k = 0,\ldots, n) \qquad (8)$$

$$x_{jk} - x_{j,k+1} \leq u_{jk} \qquad (j \in W, k = 1,\ldots, n-1) \qquad (9)$$

$$\sum_{j \in W} u_{jk} \leq \alpha_k \qquad (k = 1,\ldots, n-1) \qquad (10)$$

$$x_{jk} \in \{0,1\} \qquad (j \in W, k = 0,\ldots, n) \qquad (11)$$

$$y_{ik} \in \{0,1\} \qquad (i \in V, k = 0,\ldots, n) \qquad (12)$$

$$u_{jk} \in \{0,1\} \qquad (j \in W, k = 1,\ldots, n-1). \qquad (13)$$

In this model, constraints (7) and (8) play the same role as constraints (2) and (3) in MCLP-n. Constraints (9), (10) and (13) control the number of waiting site changes when the system moves from state k to state $k + 1$. It is easy to see that MECRP reduces to MCLP-n when $q_n = 1$. The relaxation procedure we have proposed for the MECRP solves this problem optimally whenever $\alpha_k \geq k$ for $k = 1,\ldots, n - 1$. Note that our formulation does not contain relocation costs, but only upper bounds on the number of waiting site changes. Imposing this rule makes our model tractable although it may adversely affect the expected coverage. However, limiting the number of vehicle relocations is important in practice for physicians. The same requirement was present in an ambulance relocation problem recently studied by the authors[17]. An alternative rule could

have been to impose an upper limit on the relocation distance but because the number of vehicles is rather small, this would have resulted in an even worse expected coverage.

Our model assures there is no direct cost in implementing a new solution when the system moves to a different state. The only information required to maximize expected covered demand is the proportion of time the system is in each state, as opposed to state counts which would have been needed if transition costs had been used. To apply MECRP in a dynamic context it suffices to solve this problem only once and to implement the location decisions given by the variables x_{jk} whenever the system is in state k : one of the vehicles is assigned to a call and the relocation of the remaining $k-1$ vehicles from their current to their new waiting sites is achieved by solving a transportation problem.

Computational results

The integer program (6)-(13) was run on CPLEX 7.0 on a Sun Ultra 5 Station (400 MHz). We first executed the model using the following data corresponding to the neighbouring cities of Montreal and Laval with $|V| = 3049$, $|W| = 64$, and $n = 3$, 4, 5 and 6. The demand points and the waiting sites are shown in Figure 1. Each of the 3049 demand points $i \in V$ is the result of an aggregation process and has an associated demand d_i ranging between 2 and 7000. Travel times were computed using a Manhattan metric and a constant vehicle speed of 60 km/h. We used a coverage radius r equal to 8 minutes. To compute p we used $\lambda = 1.03$ and $\mu = 1.46$. These figures were obtained from real data provided by Urgences Santé which runs the emergency medical services in Montreal and Laval. In the data set provided, 869 calls were observed over a five week period, yielding an average hourly arrival rate of 1.03. The average service time was 41 minutes, *i.e.*, $60 / \mu$.

The model was solved for the following values of α_k controlling the number of waiting site changes : $\alpha_k = 0$, $\alpha_k = 1$, $\alpha_k = \lceil k/2 \rceil$, and $\alpha_k = k$. Figure 2 and Table 1 provide the waiting locations suggested by the model for $n = 3$, 4, 5 and 6, and $k = 1, \ldots, n$. Expected coverage ratios equal to $\sum_{k=1}^{n} \sum_{i \in V} d_i q_k y_{ik} \; / \; \sum_{i \in V} d_i$ and computation times are provided in Table 2.

Table 1. Waiting sites for several values of n and k

n	k	$\alpha_k = 0$	$\alpha_k = 1$	$\alpha_k = \lceil k/2 \rceil$	$\alpha_k = k$
3	1	5	5	5	5
	2	5-9	5-9	5-9	5-9
	3	5-9-20	9-31-35	9-31-35	9-31-35
4	1	5	5	5	5
	2	5-9	5-9	5-9	5-9
	3	5-9-20	9-20-28	9-31-35	9-31-35
	4	5-9-12-20	9-20-23-39	9-20-23-39	9-20-23-39
5	1	5	5	5	5
	2	5-9	5-9	5-9	5-9
	3	5-9-20	9-20-28	9-31-35	9-31-35
	4	5-9-12-20	9-20-23-39	9-20-23-39	9-20-23-39
	5	5-9-12-20-55	1-9-12-20-23	1-9-12-20-23	1-9-12-20-23
6	1	28	5	5	5
	2	9-28	5-9	5-9	5-9
	3	8-9-28	9-20-28	9-31-35	9-31-35
	4	8-9-28-33	9-20-23-39	9-20-23-39	9-20-23-39
	5	8-9-12-28-33	1-9-12-20-23	1-9-12-20-23	1-9-12-20-23
	6	6-8-9-12-28-33	1-6-9-12-20-23	1-6-9-12-20-23	1-6-9-12-20-23

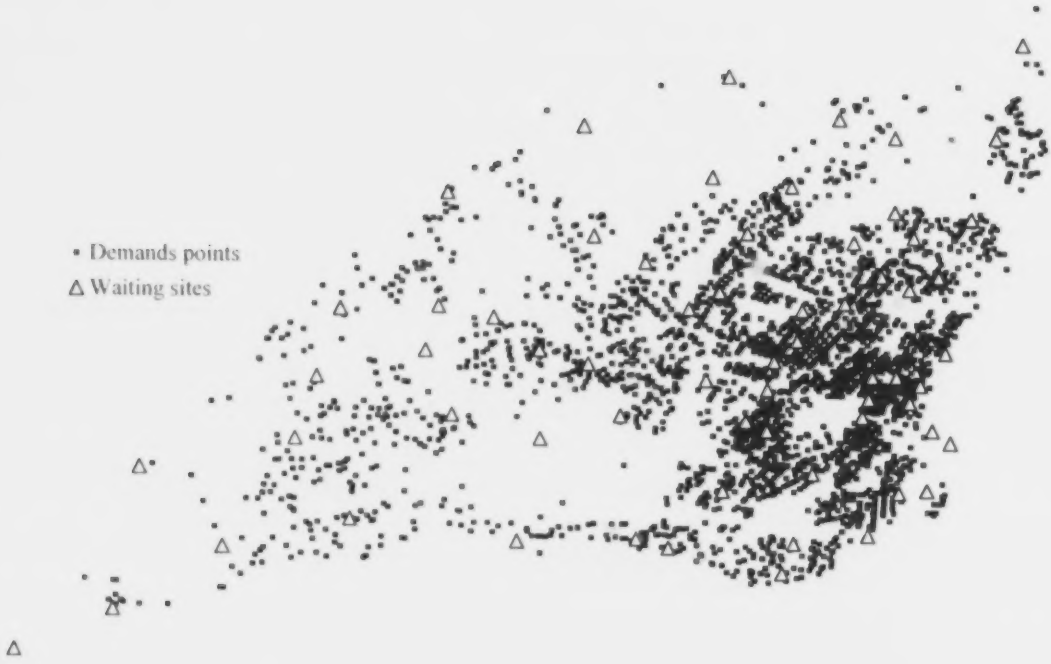

Figure 1. Demand points and waiting sites in Montreal and Laval

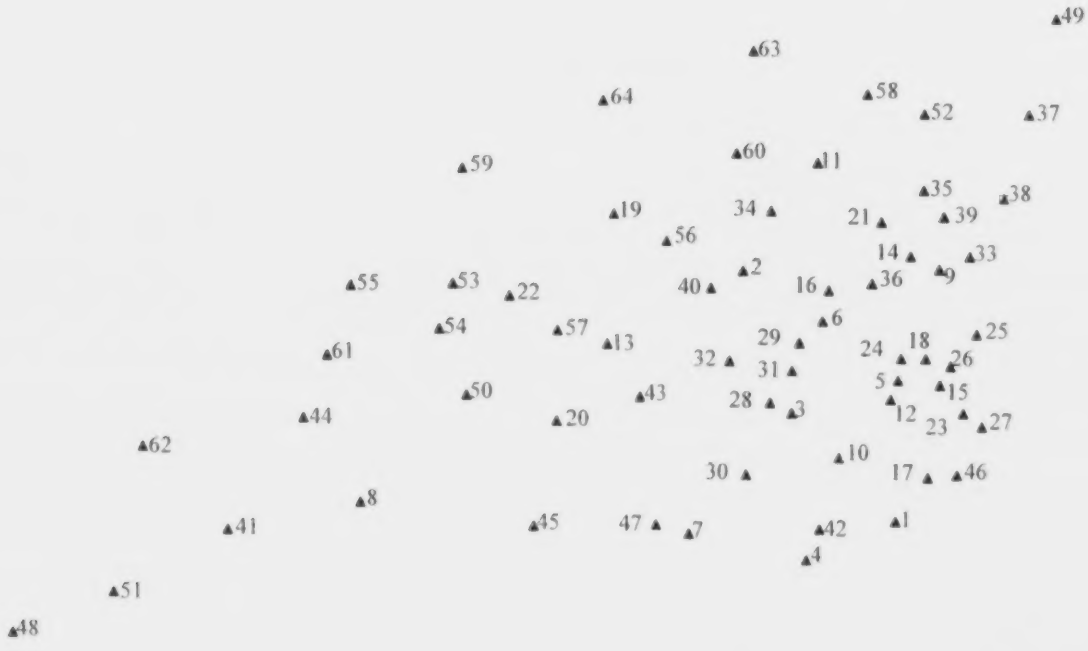

Figure 2. Waiting sites indices used in Table 1.

As expected, the column $\alpha_k = 0$ of Table 1 displays embedded waiting site sets since all sites present at state k are also used at state $k + 1$ ($k \leq n - 1$). In addition, the expected coverage ratio increases as α_k becomes larger but no difference is observed when going from $\alpha_k = \lceil k/2 \rceil$ to $\alpha_k = k$; in other words no added benefit is obtained by letting more than half of all waiting sites to change. With one exception ($n = 6$ and $\alpha_k = 0$), sites number 5 and 9 are always selected. This makes sense since these locations are located in one of the most densely populated area of Montreal (see Figure 1). When n increases more excentric sites are selected. The choice of site 28 when $n = 6$ and $\alpha_k = 0$ is explained by the multiplicity of equivalent optimal solutions for this case.

Table 2. Expected coverage ratios (first line) and computation times in minutes (second line)

n	$\alpha_k = 0$	$\alpha_k = 1$	$\alpha_k = \lceil k/2 \rceil$	$\alpha_k = k$
3	0.377	0.380	0.380	0.380
	0.6	0.5	0.5	0.5
4	0.483	0.487	0.489	0.489
	1.2	9.4	0.8	0.9
5	0.564	0.574	0.576	0.576
	111.1	27.1	21.8	7.2
6	0.637	0.643	0.643	0.643
	194.8	78.8	29.6	19.3

Table 2 also indicates that depending on n and α_k, computation times vary between less than one minute for $n = 3$ to slightly more than three hours for $n = 6$. These rather important computing times are due to the very large number of binary variables in the MECRP model. However, one should bear in mind that this model is used as a medium term planning tool (*e.g.*, a season) and is not frequently solved.

Simulations were then performed using the results of the model for $n = 3, 4, 5, 6$ and the actual calls received by Urgences Santé over the five week period. We distinguish between the "static

case" and "$\alpha_k = 0$". In the static case no vehicle is relocated. For example, if there are three waiting sites 5, 9 and 20 and the vehicle located at site 5 answer a call, then the remaining two vehicles do not move. Under the policy $\alpha_k = 0$, the waiting sites do not change but some vehicles may move. Using the same example, the vehicle located at 20 moves to site 5. In other words, waiting sites are modified as prescribed by the solution whenever a call occurs and a transportation problem is solved to determine vehicle relocations. The vehicle dispatched to answer a call is always the closest one. The following statistics on the system performance were compiled :

1) average coverage ratio;
2) average response time (minutes);
3) maximum response time (minutes);
4) average repositioning distance (kilometers);
5) maximum repositioning distance (kilometers);
6) number of calls delayed out of 869 (these are calls that could not be immediately served because no physician car was available);
7) proportion of calls covered within eight minutes.

Table 3. Summary of simulation results

n	Performance statistics	Static case	$\alpha_k = 0$	$\alpha_k = 1$	$\alpha_k = \lceil k/2 \rceil$	$\alpha_k = k$
3	Average coverage ratio	-	0.389	0.393	0.393	0.393
	Average response time	12.02	11.04	10.78	10.78	10.78
	Maximum response time	56.39	52.79	52.79	52.79	52.79
	Average repositioning distance	-	8.86	5.76	5.76	5.76
	Maximum repositioning distance	-	15.28	9.66	9.66	9.66
	Number of calls delayed	37	34	33	33	33
	Prop. of calls covered within 8 min.	0.392	0.426	0.457	0.457	0.457
4	Average coverage ratio	-	0.508	0.514	0.516	0.516
	Average response time	9.69	9.16	9.10	8.95	8.95
	Maximum response time	49.40	49.40	52.79	52.79	52.79
	Average repositioning distance	-	9.41	5.83	4.90	4.90
	Maximum repositioning distance	-	17.54	10.32	9.67	9.67
	Number of calls delayed	5	3	3	3	3
	Prop. of calls covered within 8 min.	0.510	0.537	0.557	0.586	0.586
5	Average coverage ratio	-	0.594	0.607	0.608	0.608
	Average response time	9.14	8.33	7.97	7.89	7.89
	Maximum response time	49.40	49.40	49.40	49.40	49.40
	Average repositioning distance	-	10.32	5.05	4.85	4.85
	Maximum repositioning distance	-	19.01	15.16	15.16	15.16
	Number of calls delayed	1	0	0	0	0
	Prop. of calls covered within 8 min.	0.557	0.606	0.678	0.684	0.684
6	Average coverage ratio	-	0.672	0.677	0.676	0.676
	Average response time	7.82	7.56	7.09	7.09	7.09
	Maximum response time	37.72	37.72	37.72	37.72	37.72
	Average repositioning distance	-	11.60	7.49	6.99	6.99
	Maximum repositioning distance	-	24.71	25.33	25.33	25.33
	Number of calls delayed	0	0	0	0	0
	Prop. of calls covered within 8 min.	0.652	0.693	0.749	0.748	0.748

Simulation results reported in Table 3 indicate that applying a relocation policy is better than using a static approach. Thus, in our simulations, moving from the static case to $\alpha_k = 0$ decreases the average response time between 3.33% (for $n = 6$) and 8.88% (for $n = 5$); similarly the proportion of calls covered within eight minutes increases from 5.29% (for $n = 4$) to 8.79% (for $n = 5$). Allowing even more relocations produces more dramatic results; when $\alpha_k = \lceil k/2 \rceil$ the average response time is reduced between 7.63% (for $n = 4$) and 13.67% (for $n = 5$); similarly the proportion of calls covered within 8 minutes increases from 14.72% (for $n = 6$) to 22.80% (for $n = 5$). Note, however, that allowing up to k waiting site relocations does not provide any improvement over the case where $\lceil k/2 \rceil$ relocations are allowed. One interesting observation is that using n vehicles and allowing all waiting sites to change ($\alpha_k = k$) can have about the same effect as using $n + 1$ vehicles and allowing no vehicle relocation. Thus using five vehicles with $\alpha_k = k$ yields a better proportion of calls covered within eight minutes (0.684) than using six vehicles without vehicle relocations (0.652). The average response time (7.89 minutes) is slightly larger in the first case than in the second case (7.82 minutes). These observations must be interpreted with care since they do not hold for all values of n. It is interesting to see that the average coverage ratios observed in the simulations are consistent with the theoretical values of Table 2, which indicates that our model matches reality.

As expected, when α_k increases, the average and maximum repositioning distances decrease. This is not directly controlled by the model but results from the fact that allowing more vehicles to relocate will result in a shorter relocation distance for each of them. This phenomenon was also observed in a previous ambulance relocation study[17]. Finally, the average and maximum response times also decrease when α_k increases. This is an indirect consequence of our model because allowing more waiting site relocations also increases the proportion of the demand covered within a preset time standard and hence decreases the average response time.

Conclusions

We have formulated and solved a dynamic relocation problem arising in the management of emergency vehicles. The aim was to provide a relocation policy ensuring an adequate expected

demand coverage while controlling the number of relocations. In contexts where the number of vehicles is relatively small, as is the case of physicians cars, the problem can be modeled and solved directly using a standard optimizer like CPLEX. Simulations conducted using a five week sample of calls in Montreal and Laval confirm the feasibility of the proposed approach. Our results clearly demonstrate the benefits of relocating emergency vehicles in terms of average response time and proportion of calls covered within eight minutes. To our knowledge, very few studies have addressed dynamic emergency vehicle relocation problems. Contrary to what was done by Gendreau, Laporte and Semet[17] and to the current practice in dynamic vehicle routing[20], the dynamic problem we have studied is rather interesting in that it can be solved optimally through the execution of a single integer linear program.

Acknowledgements

This work was partly supported by the Fonds de la recherche en santé du Québec under grant 980861, by the Canadian Natural Sciences and Engineering Council under grants OGP0038816, OGP0039682 and by the Région Nord–Pas de Calais, France. This support is gratefully acknowledged. Thanks are also due to Émilie Frot and Geneviève Hernu for their help with programming and to Marko Blais from Urgences Santé who provided some data. We are also grateful to the referees for their valuable comments.

References

1. C. Toregas, R. Swain, C.S. ReVelle and L. Bergman (1971) The location of emergency service facilities. *Operations Research* 19, 1363-1373.

2. R.L. Church and C.S. ReVelle (1974) The maximal covering location problem. *Papers of the Regional Science Association* 32, 101-118.

3. D.A. Schilling, D.J. Elzinga, J. Cohon, R.L. Chuch and C.S. ReVelle (1979) The TEAM/FLEET models for simultaneous facility and equipment sitting. *Transportation Science* 13, 163-175.

4. M.S. Daskin and E.H. Stern (1981) A hierarchical objective set covering model for emergency medical service vehicle deployment. *Transportation Science* 15, 137-152.

5. K. Hogan and C.S. ReVelle (1986) Concept and applications of backup coverage. *Management Science* 32, 1434-1444.

6. M. Gendreau, G. Laporte and F. Semet (1997) Solving an ambulance location model by tabu search. *Location Science* 5, 75-88.

7. R. C. Larson (1972) *Urban Police Patrol Analysis*. The MIT Press, Cambridge, MA.

8. R. C. Larson (1974) A hypercube queuing model for facility location and redistricting in urban emergency services. *Computers & Operations Research* 1, 67-75.

9. R. C. Larson (1975) Approximating the performance of emergency service systems. *Operations Research* 23, 845-868.

10. M.S. Daskin (1983) The maximal expected covering location model: formulation, properties, and heuristic solution. *Transportation Science* 17, 48-70.

11. J.F. Repede and J.J. Bernardo (1994) Developing and validating a decision support system for locating emergency medical vehicles in Louisville, Kentucky. *European Journal of Operational Research* 75, 567-581.

12. R. Batta, J.M. Dolan, N.N. Krishnamorthy (1989) The maximal expected covering location model revisited. *Transportation Science* 23, 277-287.

13. V. Marianov and C.S. ReVelle (1994) The queuing probabilistic location set covering problem and some extensions. *Socio-Economic Planning Sciences* 28, 167-178.

14. C.S. ReVelle and K. Hogan (1989) The maximum availability location problem. *Transportation Science* 23, 192-200.

15. M.O. Ball and F.L. Lin (1993) A reliability model applied to emergency service vehicle location. *Operations Research* 41, 18-36.

16. P. Kolesar and W.E. Walker (1974) An algorithm for the dynamic relocation of fire companies. *Operations Research* 22, 249-274.

17. M. Gendreau, G. Laporte and F. Semet (2001) A dynamic model and parallel tabu search heuristic for real-time ambulance relocation. *Parallel Computing* 27, 1641-1653.

18. V. Marianov and C.S. ReVelle (1995) Siting emergency services. In: *Facility Location : A Survey of Applications and Methods*, Z. Drezner (ed.) Springer-Verlag, New York, 199-223.

19. L. Brotcorne, G. Laporte and F. Semet (2003) Ambulance location and relocation models. *European Journal of Operational Research* 147, 451-463.

20. M. Gendreau, F. Guertin, J.-Y. Potvin and É. Taillard (1999) Parallel tabu search for real-time vehicle routing and dispatching. *Transportation Science* 33, 381-390.